

AD-A052 445

MASSACHUSETTS INST OF TECH CAMBRIDGE ARTIFICIAL INTE--ETC F/6 9/2
THE INCREMENTAL GARBAGE COLLECTION OF PROCESSES.(U)

DEC 77 H G BAKER, C HEWITT

N00014-75-C-0643

UNCLASSIFIED

AI-M-454

NL

| OF |
ADA
052445

END
DATE
FILMED
5 - 78

DDC

ADA052445

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

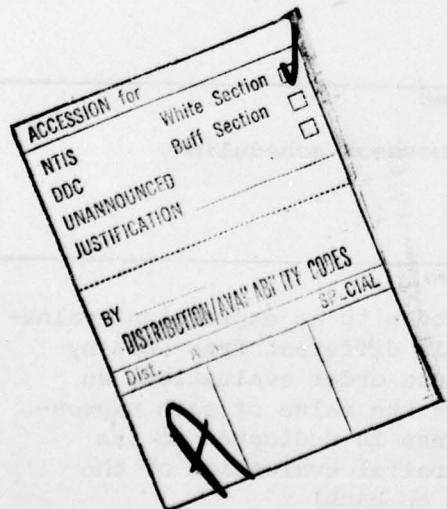
REPORT DOCUMENTATION PAGE			READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER ⑯ ATM-454	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER ⑯	
4. TITLE (and subtitle) The Incremental Garbage Collection of Processes	5. TYPE OF REPORT & PERIOD COVERED MEMORANDUM rept.,		
6. AUTHOR(s) ⑯ Henry G. Baker, Jr. & Carl Hewitt	7. CONTRACT OR GRANT NUMBER(s) N00014-75-C-0643 N00014-75-C-0522		
8. PERFORMING ORGANIZATION NAME AND ADDRESS Artificial Intelligence Laboratory 545 Technology Square Cambridge, Massachusetts 02139	9. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS		
10. CONTROLLING OFFICE NAME AND ADDRESS Advanced Research Projects Agency 1400 Wilson Blvd Arlington, Virginia 22209	11. REPORT DATE ⑯ December 1977		
12. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Office of Naval Research Information Systems Arlington, Virginia 22217	13. NUMBER OF PAGES 10		
14. SECURITY CLASS. (of this report) UNCLASSIFIED	15. DECLASSIFICATION/DOWNGRADING SCHEDULE		
16. DISTRIBUTION STATEMENT (of this Report) Distribution of this document is unlimited.	17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report) D D C REF ID: A APR 10 1978 F		
18. SUPPLEMENTARY NOTES None			
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) garbage collection, multiprocessing systems, processor scheduling, "lazy" evaluation, "eager" evaluation			
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) This paper investigates some problems associated with an expression evaluation order that we call "future" order, which is different from call-by-name, call-by-value, and call-by-need. In future order evaluation, an object called a "future" is created to serve as the value of each expression that is to be evaluated and separate process is dedicated to its evaluation. This mechanism allows the fully parallel evaluation of the expressions in a programming language. (cont'd on back)			

BLOCK 20 CONCLUDED

We discuss an approach to a problem that arises in this context:futures which were thought to be relevant when they were created become irrelevant through not being needed later in the computation. The problem of irrelevant processes also appears in multiprocessing problem-solving systems which start several processors working on the same problem but with different methods, and return with the solution which finishes first. This parallel method strategy has the drawback that the processes which are investigating the losing methods must be identified, cleanly stopped and the processors they are using re-assigned to more useful tasks.

The solution we propose is that of incremental garbage collection. The goal structure of the solution plan should be explicitly represented in memory as part of the graph memory (like LISP's heap) so that a garbage collection algorithm can discover which processes are performing useful work, and which can be recycled for a new task.

An incremental algorithm for the unified garbage collection of storage and processes is described.



MASSACHUSETTS INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY

AI Memo 454

December 1977

The Incremental Garbage Collection of Processes

Henry G. Baker, Jr. and Carl Hewitt

This paper investigates some problems associated with an expression evaluation order that we call "future" order, which is different from call-by-name, call-by-value, and call-by-need. In future order evaluation, an object called a "future" is created to serve as the value of each expression that is to be evaluated and separate process is dedicated to its evaluation. This mechanism allows the fully parallel evaluation of the expressions in a programming language.

We discuss an approach to a problem that arises in this context: futures which were thought to be relevant when they were created become irrelevant through not being needed later in the computation. The problem of irrelevant processes also appears in multiprocessing problem-solving systems which start several processors working on the same problem but with different methods, and return with the solution which finishes first. This *parallel method* strategy has the drawback that the processes which are investigating the losing methods must be identified, cleanly stopped, and the processors they are using re-assigned to more useful tasks.

The solution we propose is that of incremental garbage collection. The goal structure of the solution plan should be explicitly represented in memory as part of the graph memory (like Lisp's heap) so that a garbage collection algorithm can discover which processes are performing useful work, and which can be recycled for a new task.

An incremental algorithm for the unified garbage collection of storage and processes is described.

Key Words and Phrases: garbage collection, multiprocessing systems, processor scheduling, "lazy" evaluation, "eager" evaluation.

CR Categories: 3.60, 3.80, 4.13, 4.22, 4.32.

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the laboratory's artificial intelligence research is provided in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-75-C-0522.

1. Introduction

Processors are becoming very cheap and there is good evidence that this trend will continue in the next few years. As a result, there has been considerable interest in how to apply large numbers of processors to the solution of a single task [8]. Since efficient utilization of a horde of processors requires a lot of communication, sorting networks have been devised [2,21] which allow every processor in an N -processor system to both send and receive a message on every clock pulse. Furthermore, the transit time through the network is on the average only $O(\log N)$ and the size of the network only $O(N \log^2 N)$. However, it is still not clear how to effectively utilize all of these processors.

Other researchers [11,5,22] quite rightly note that languages without side-effects, e.g. "pure" LISP, are excellently suited for the purpose of representing many algorithms intended for execution on a host of processors since their lack of side-effects eliminates a great source of complexity in parallel execution. Thus, "Church-Rosser" theorems can be proved which ensure the invariance of the value of an expression regardless of the order or relative speed of evaluation. However, we must keep in mind that this kind of parallelism does not implement the most general form of communication between processes. For example, an airline reservation system cannot be implemented in such a language, due to its non-determinate behavior.

In this paper, we consider an "eager beaver" evaluator for a language without side-effects, such as pure LISP. When an expression of the language is given to the evaluator by the user, the evaluator evaluates it and all of its subexpressions as soon as possible, and in parallel. The evaluator does this by creating and returning for each subexpression a *future*, which is a promise to deliver the value of that subexpression at some later time, if it has a value. Each future can evaluate its subexpression independently and concurrently with other futures because it is created with its own evaluator *process*, which is dedicated to evaluating its subexpression. When the value of a future is needed explicitly, e.g. by the primitive function "+", the evaluator of the subexpression may or may not be complete. If it is complete, the future's value is immediately available; if not, the requesting process is forced to wait until the evaluation of the subexpression is done.

Futures are created recursively in the evaluation of an expression by our eager evaluator whenever it encounters functional application. A new future is created for each argument, resulting in the parallel (collateral) evaluation of those arguments, while the main process tackles the job of evaluating the function position and applying it to the tuple of argument futures. We call the main evaluator process the *parent*, while the futures it directly creates are its *offspring*.

More precisely, a *future* is a triple (*process*, *cell*, *waiting room*), where *process* is the virtual processor initialized to evaluate an argument expression in its proper environment, *cell* is a writable location in memory which will save the value of the argument when it is ready, to avoid recomputing it, and *waiting room* is a set of processes which are waiting for the value of this future.

When the future is created, its *process* starts evaluating the subexpression in the given environment. If any other process needs the value of this future before it is ready, the requesting process enters the *waiting room* of the future and goes to sleep. When the value promised by the future is ready, its *process* stores that value into the future's *cell*, wakes up all of the processes in the future's *waiting room*, and dies. Henceforth, any process needing this future's value can find it

in the future's cell, without waiting or performing any further computation.

Notice that a *future* of an expression is different from the delayed value [26,23,12,15,10] of the expression in that the latter are designed to delay evaluation of the expression until the value is needed while a future immediately dedicates a processor to evaluating the expression. This difference is both a strength and a weakness of eager evaluation.

The main problem with our eager interpreter is that it can be wasteful, because it anticipates which values are going to be required to compute the final result. For example, a process may be assigned to the computation of a future whose value will never be needed; in this case, we say that the process is *irrelevant*. If there were no way to determine irrelevancy, these irrelevant processes could tie up a significant amount of computing power. Furthermore, if a process was assigned to evaluate a non-terminating expression, its computational power would be lost to the system forever! In the following sections, we argue that the "garbage collection" of passive storage can be extended to the reclamation of these irrelevant active processes. Furthermore, we show that this garbage collection can be done *incrementally*, thus eliminating the long delays classically associated with garbage collection.

2. Garbage Collecting Irrelevant Futures

A classical garbage collector for passive storage starts by marking the *root* of the heap of passive storage nodes, and proceeds by propagating marks from marked nodes to their offspring, until there is no unmarked node with a marked parent. Upon the completion of this process, any nodes which are still unmarked are not accessible from the root; hence they are declared garbage and returned to the list of available free nodes.

The key to garbage collecting *active processes* is that active process's process-states are addressable as vectors of words in the common address space of all the processors, but marked with a special type code. These vectors store the contents of the registers of the processes, plus additional status information. We claim that processes whose process-states become inaccessible from the root are irrelevant and should be reclaimed. The top-level process--that assigned to the top level future--is always relevant since the user expects an answer, and therefore it is always directly accessible from the root of the heap. Any offspring of this future whose values are still required are accessible to it. Hence by induction, relevant processes remain accessible from the root. If a future becomes inaccessible from the root, then no other process can access its value--even when it is ready--and hence the future and its process are irrelevant.

In order that all irrelevant processes be identified as soon as possible, we must make sure that all processes classified as accessible are truly relevant to the computation. An example of a process which is accessible but irrelevant is that of a *loiterer*, i.e. a process which is accessible only through the "waiting room" of some future. A loiterer is waiting for the value of one or more futures, but whose value is not needed by any other process. Loiterers cannot be immediately garbage collected because of the outstanding waiting-room pointers to them, but they need not be restarted. If a loiterer is encountered in a second garbage collection, it can at that time be reclaimed. Hence, waiting-room accessibility is a second-class form of accessibility which will not protect a loiterer from eventually being garbage collected.

If busy waiting is used, waiting rooms are not necessary, and thus their pointers do not require special treatment by the garbage collector. However, busy waiting requires that a high price be paid for communication channels between the waiter and the waitee, because the incessant queries clog these channels.

Garbage collection is made incremental by using some of the ideas from an earlier paper [1], which in turn is based on the work of Dijkstra [6,7] and Lamport [17,18]. The mark phase of our incremental garbage collector process employs three colors for every object--white, grey, and black. Intuitively, *white* nodes are not yet known to be accessible, *grey* nodes are known to be accessible, but whose offspring have not yet been checked, and *black* nodes are accessible, and have accessible offspring. Initially, all nodes (including processes) are white. A white node is made grey by *shading* it; i.e. making it "at least grey" [6], while a grey node is *marked* by shading its offspring and making the node black--both indivisible processes. Marking is initiated by stopping all processes and shading the root. Marking proceeds by finding a grey node, shading its offspring, then making that node black. When there are no more grey nodes, garbage collection is done; all still-white nodes are then emancipated and the colors white and black switch interpretations.

Although all user processes must be stopped when garbage collection is begun, a user process can be restarted as soon as it has been blackened by the collector. Since the top-level process is pointed at directly by the root of the heap, it is restarted almost immediately. It should be obvious that when a process first becomes black, it cannot point directly at a white node. We wish to preserve this assertion. Therefore, whenever a running black process is about to violate it--by inserting into one of its registers the white component of a node it is already pointing at--it immediately shades the white node before proceeding. Furthermore, every new node the process needs is created black. The intuitive rationale behind these policies is that *so far as any black process is concerned, the garbage collection has already finished*. Furthermore, *the nodes which are found accessible by the garbage collector are exactly those which were accessible at the time the garbage collection was started*.

We prove the correctness of this garbage collector informally. The garbage collector is given a head start on all of the processes because they are stopped when it is started. When a process is restarted, it is black, and everything it sees is at least grey, hence it is in the collector's wake. Whenever a process attempts to catch up to the collector by tracing an edge from a node it can access directly, that node is immediately shaded. Therefore, it can never pass or even catch the collector. Since the collector has already traced any node a process can get its registers on, the process cannot affect the connectivity of the nodes that the collector sees. Because white or grey processes are not allowed to run, any created nodes are black, and since nodes darken monotonically, the number of white nodes must monotonically decrease, proving termination.

Our garbage collector has only one phase--the mark phase--because it uses a compacting, copying algorithm [9,4] which marks and copies in one operation. This algorithm copies accessible list structures from an "old semispace" into a "new semispace". As each node is copied, a "forwarding address" is left at its old address in the old semispace. If the Minsky copying algorithm is used [9], the collector has its own stack to keep track of grey nodes; the Cheney algorithm [4] uses a "scan pointer" to linearly scan the new semispace, while updating the pointers of newly moved nodes by moving the nodes they point to. The correspondence between our coloring scheme and these algorithms is this: *white* nodes are those which reside in the old semispace; *grey* nodes are those which have been copied to the new semispace, but whose outgoing pointers have not been updated to point into the new semispace (i.e. have not yet been encountered by the scan pointer in the Cheney algorithm); and *black* nodes are those which have been both moved and updated (i.e. are

behind the scan pointer). When scanning is done (i.e. there are no more grey nodes and all accessible nodes have been copied), the old and new semispaces then interchange roles. Reallocating processors is simple; all processors are withdrawn at the start of garbage collection, and are allocated to each process as it is blackened. Thus, when the garbage collection has finished, all and only relevant processes have been restarted.

The restriction that white or grey processes cannot run can be relaxed to allow white processes to run so long as a white process does not cause a black node to point to a white one. This can only happen if the white process is trying to perform a side-effect (RPLACA or RPLACD in LISP) on a black node. If operations of this type are suspended until either the process either becomes black or is garbage-collected, then proper garbage collector operation can be ensured, and convergence guaranteed. Under these conditions, a process creates new cells of its own color, i.e. white processes create only white cells. When a white process is encountered by the garbage collector, it must stop and allow itself to be colored black before continuing.

The notion that processes must be marked as well as storage may explain some of the trouble that Dijkstra and Lamport had when trying to prove their parallel garbage collection algorithm correct [6,7,17,18]. Since their algorithm does not mark a user process by coloring it black (thereby prohibiting it from directly touching white nodes), and allows these white processes to run, the proof that the algorithm collects only and all garbage is long and very subtle (see [18]).

3. Coroutines and Generators

One problem with our "eager beaver" evaluator is that some expressions which have no finite values will continue to be evaluated without mercy. Consider, for example, the infinite sequence of squares of integers 0,1,4,9,... We give below a set of LISP-like functions for computing such a list.

```

squares-beginning-with ≡ (λx. (cons (* x x) (squares-beginning-with (+ x 1)))) ; Compute an element.

cons ≡ (λx y.
          (λmessage.
            (if (= message 'car) x
                (= message 'cdr) y))) ; Define CONS function.

car ≡ (λx. (x 'car)) ; Ask for first component.

cdr ≡ (λx. (x 'cdr)) ; Ask for second component.

list-of-squares ≡ (squares-beginning-with 0) ; Start the recursion.

```

The evaluation of "(squares-beginning-with 0)" will start off a future evaluating "(cons ...)", which will start up another future evaluating "(squares-beginning-with 1)", and so forth. Since this computation will not terminate, we might worry whether anything useful will ever get done. One way to ensure that this computation will not clog the system is to convert it into a "lazy" computation [26,23,12,10] by only allowing it to proceed past a point in the infinite list when someone forces it to go that far. This can be easily done by performing a lambda abstraction on the expression whose evaluation is to be delayed. Since our evaluator will not try to further evaluate a λ-expression, this will protect its body from evaluation by our eager beavers.

```
squares-beginning-with' ≡
(λx. (cons ⟨* x x⟩
          (λmessage. ((squares-beginning-with' (+ x 1)) message)))) ; Protect from early evaluation.
```

However, this technique is not really necessary if we use an *exponential scheduler* for the proportion of effort assigned to each process. This scheduler operates recursively by assigning 100% of the system effort to the top-level future, and whenever this future spawns new futures, it allocates only 50% of its allowed effort to its offspring. While a process is in the waiting room of a future, it lends its processing effort to the computation of that future. However, a future which finishes returns its effort to helping the system--not its siblings. Now the set of futures can be ordered according to who created whom and this ordering forms a tree. As a result of our exponential scheduling, the further down in this tree a future is from the top-level future, the lower its priority in scheduling. Therefore, as our eager beavers produce more squares, they become exponentially more discouraged. But if other processes enter the waiting room for the square of a large number, they lend their encouragement to its computation.

Call-by-future evaluation provides for the maximal concurrency possible in evaluating the expressions of a language. It can provide more parallelism than current data flow machines [0,5] or "eventual values" [16]. For example consider the following program which computes the square root of the sum of the squares of its arguments

$$f \equiv (\lambda x y. (\text{square-root} (+ (* x x) (* y y))))$$

Note that in computing the value of an expression such as the following ($k 3 (f (h 3) (g 4))$) that the square of ($h 3$) can be performed in parallel with the square of ($g 4$). In addition the square root of the sum of these values might be performed after the function k has been entered! Thus there is a great deal of potential concurrency in the evaluation of the above expression.

In an evaluator which uses call-by-future for CONS, the obvious program for MAPCAR (the LISP analog of APL's parallel application of a function to a vector of arguments) will automatically do all of the function applications in parallel in a "pipe-lined" fashion. However, due to the scheduler the values earlier in the list will be accorded more effort than the later ones.

Because this scheduler is not omniscient, system effort will still have to be reallocated by the garbage collector as it discovers irrelevant processes and returns their computing power to help with still relevant tasks.

4. Time and Space

"Lazy" evaluation [26,12,10] is an optimal strategy [23,3] for evaluating expressions on a single processor, in the sense that the minimum number of reductions (procedure calls) are made. However, when more than one processor is available to evaluate the expression, it is not clear what strategy would be optimal. If nothing is known about the particular expression being evaluated, we conjecture that any reasonable strategy must allocate one processor to lazy evaluation, with the other processors performing eager evaluation. We believe that our "eager beaver" evaluator implements this policy, and unless the processors interfere with one another excessively, a computation must always run faster with an eager evaluator running on multiple processors than a lazy evaluator running on a single processor. If there are not enough processors to allocate one for every future, then we believe that our "exponential scheduling" policy will do a good job of

dynamically allocating processor effort where it is most needed.

Although the universal creation of futures should reduce the time necessary to evaluate an expression, we must consider how the space requirements of this method compare with other methods. The space requirements of futures are hard to calculate because under certain schedules, future order evaluation approximates call-by-value, while with other schedules, it is equivalent to call-by-name (but evaluated only once). In the worst case, the space requirements of futures can be arbitrarily bad, depending upon the relative speed of the processors assigned to non-terminating futures.

5. The Power of Futures

The intuitive semantics associated with a future is that it runs asynchronously with its parent's evaluation. This effect can be achieved by either assigning a different processor to each future, or by multiplexing all of the futures on a few processors. Given one such implementation, the language can easily be extended McCarthy [19] with a construct having the following form: "(EITHER e_1 e_2 ... e_n)" means evaluate the expressions e_i in parallel and return the value of "the first one that finishes". Ward [24] shows how to give a Scott-type lattice semantics for this construct. He starts with a power-set of the base domain and gives it the usual subset lattice structure, then extends each primitive function to operate on sets of elements from the base domain in the obvious way, and finally defines the result of the EITHER construct to be the *least upper bound (LUB)* of all the e_i in the subset lattice. The EITHER construct is approximated¹ by spawning futures for all the e_i , and polling them with the parent process until the first one finishes. At that point, its answer is returned as the value of the "EITHER" expression, and the other futures become inaccessible from the root of the heap.

We give several examples of the power of the "EITHER" construct:

```
(multiply x y) ≡ (EITHER (if x=0 then 0 else (loop))
                           (if y=0 then 0 else (loop)))
                           (* x y))
```

```
(integrate exp bound-variable) ≡
  (EITHER (fast-heuristic-integrate exp bound-variable)
           (Risch-integrate exp bound-variable))
```

The first example is that of a numeric product routine whose value is zero if either of its arguments are zero, even if the non-zero argument is undefined. The second example is an integration routine for use in a symbolic manipulation language like Macsyma, where there is a relatively fast heuristic integration routine which looks for common special cases, and a general but slow decision procedure called the Risch algorithm. Since the values of both methods are guaranteed to be the same (assuming that they perform integration properly), we need not worry about the possibility of non-determinacy of the value of this expression (i.e. non-singleton subsets of the base domain in Ward's lattice model).

1. This implementation is only an approximation because only singleton sets of elements of the base domain can ever be returned.

One may ask what the power of such an "EITHER" construct is; i.e. does it increase the expressive power of the language in which it is embedded? A partial answer to this question has been given with respect to "uninterpreted" schemata. Uninterpreted schemata answer questions about the expressive power of programming language constructs which are implicit in the language, rather than being simulated. For example, one can compare the power of recursion versus iteration in a context where stacks cannot be simulated. Hewitt and Paterson [13] have shown that uninterpreted "parallel" schemata are strictly more powerful than recursive schemata. The essence of this difference is that parallel schemata can simulate non-deterministic computation without bogging down in some infinite branch by following all branches in parallel.

Also, Ward [24] has shown that the "EITHER" construct strictly increases the power of the λ -calculus in the sense that there exist functions over the base domain which are inexpressible without "EITHER", but are trivially expressible with it.

6. Shared Data bases

The advantage that garbage collection has over the explicit killing of processes becomes apparent when parallel processes have access to a shared data base. These data bases are usually protected from inconsistency due to simultaneous update by a mutual exclusion method. However, if some process were to be killed while it was inside such a data base, the data base would remain locked, and hence unresponsive to the other processes requesting access.

The solution we propose is for the data base to always keep a list of pointers to the processes which it has currently inside. In this way, an otherwise irrelevant process will be accessible so long as it is inside an accessible data base. However, the moment it emerges, it will be forgotten by the data base, and subject to reclamation by garbage collection. The *crowds* component of a *serializer*, a synchronization construct designed to manage parallel access to a shared data base [14], automatically performs such bookkeeping.

7. Conclusions

We have presented a method for managing the allocation of processors as well as storage to the subcomputations of a computation in a way that tries to minimize the elapsed time required. This is done by anticipating which subcomputations will be needed and starting them running in parallel, before the results they compute are needed. Because of this anticipation, subcomputations may be started whose results are not needed, and thus our method identifies and revokes these allocations of storage and processing power through an incremental garbage collection method.

The scheme presented here assumes that all of the processors reside in a common, global address space, like that of HYDRA [25]. Since networks of local address spaces look promising for the future, methods for garbage collecting those systems need to be developed.

8. Acknowledgements

Some of the early thinking about call-by-future was done several years ago by J. Rumbaugh who was one of the first to realize that futures offer a maximum of concurrency in the execution of a program without introducing the usual pitfalls of timing errors, starvation, and deadlock. Unfortunately he did not have time to include this material in his thesis [20]. Peter Hibbard [16]

has independently discovered these virtues of futures. The main original contributions of this paper are our proposal for an exponential scheduler for "eager" evaluation and the methodology for using incremental garbage collection to reclaim irrelevant processes and redirect the scheduling priorities of processes working to produce the values of futures. Concepts similar to that of "futures" have been independently proposed by Friedman and Wise [11] and implemented by Hibbard [16].² Bob Tennent, Andrei Ershov, and Gianfranco Prini made valuable comments on this paper.

2. However, since Algol-68 does not support "returned functional values", A scheme in the language need not use garbage collection to discover irrelevant "eventuals". They can be coerced into values before being returned as the value of a procedure, and hence processor allocation can use a LIFO scheme like that used for storage of the activation records on the stack. However a certain amount of concurrency can be lost by enforcing this coercion.

9. References

- [0] Arvind and Gostelow, Kim. "Some Relationships between Asynchronous Interpreters of a Dataflow Language" IFIP Working Conference on Formal Description of Programming Concepts. 31 July to 5 August 1977.
- 1. Baker, H. G., Jr. "List Processing in Real Time on a Serial Computer". AI Working Paper 139, MIT AI Lab., Feb. 1977, also to appear in *CACM*.
- 2. Batcher, K. E. "Sorting Networks and their Applications". 1968 SJCC, April 1968, 307-314.
- 3. Berry, Gerard and Levy, Jean-Jacques. "Minimal and Optimal Computations of Recursive Programs". Record of 1977 Conference on Principles of Programming Languages, Jan. 1977, 215-226.
- 4. Cheney, C. J. "A Nonrecursive List Compacting Algorithm". *CACM* 13,11 (Nov. 1970), 677-678.
- 5. Dennis, J. and Misunas, D. P. "A Preliminary Architecture for a Basic Data-Flow Processor. In 2nd IEEE Symposium on Computer Architecture., N.Y. Jan. 1975, 126-132.
- 6. Dijkstra, E. W., Lamport, L., Martin, A. J., Scholten, C. S., Steffens, E. F. M. "On-the-fly Garbage Collection: An Exercise in Cooperation". Dijkstra note EWD496, June 1975.
- 7. Dijkstra, E. W. "After Many a Sobering Experience". Dijkstra note EWD500.
- 8. Erman, L. D. and Lesser, V. R. "A Multi-level Organization for Problem Solving using Many, Diverse, Cooperating Sources of Knowledge". IJCAI-75, Sept. 1975, 483-490.
- 9. Fenichel, R. R., and Yochelson, J. C. "A LISP Garbage-Collector for Virtual-Memory Computer Systems". *CACM* 12,11 (Nov. 1969), 611-612.
- 10. Friedman, D. P. and Wise, D. S. "CONS should not evaluate its arguments". In S. Michaelson and R. Milner (eds.), *Automata, Languages and Programming*, Edinburgh University Press, Edinburgh (1976), 257-284.
- 11. Friedman, D. P. and Wise, D. S. "The Impact of Applicative Programming on Multiprocessing". 1976 Int. Conf. on Parallel Processing, 263-272. Also *IEEE Trans. on Comps.*, to appear.
- 12. Henderson, P. and Morris J. H. "A Lazy Evaluator" In Proceedings of 3rd ACM Symposium on Principles of Programming Languages. (1976), 95-103.
- 13. Hewitt, C. and Paterson, M. "Comparative Schematology". Record of Project MAC Conference on Concurrent Systems and Parallel Computation, June 1970.
- 14. Hewitt, C. and Atkinson, R. "Parallelism and Synchronization in Actor Systems". Record of 1977 Conference on Principles of Programming Languages, Jan. 17-19, 1977, L.A., Cal., 267-280.
- 15. Hewitt, C. "Viewing Control Structures as Patterns of Passing Messages" WP 92, MIT AI Lab., Dec. 1975. Accepted for publication in the A.I. Journal.
- 16. Hibbard, P. "Parallel Processing Facilities". in *New Directions in Algorithmic Languages*, (ed.) Stephen A. Schuman, IRIA, 1976, 1-7.
- 17. Lamport, L. "Garbage Collection with Multiple Processes: An Exercise in Parallelism". Mass. Comp. Associates, CA-7602-25II, Feb. 1976.
- 18. Lamport, L. "On-the-fly Garbage Collection: Once More with Rigor". Mass. Comp. Associates, CA-7508-16II, Aug. 1975.
- 19. McCarthy, J. "A Basis of the Mathematical Theory of Computation" In P. Braffort and D. Hirschberg (eds.) *Programming Systems and Languages*. McGraw-Hill, New York(1967), 455-480.
- 20. Rumbaugh, J. E. "A Parallel Asynchronous Computer Architecture for Data Flow Programs" Ph.D. dissertation, M.I.T. May 1975. MAC TR-150.
- 21. Sullivan, H. and Bashkow T. R. "A Large Scale, Homogeneous, Fully Distributed Parallel Machine". Proc. of Fourth Annual Symposium on Computer Architecture. March 1977, 105-117.

22. Tessler, G. and Enea H. J. "A Language Design for Concurrent Processes" In Proc. SJCC, Thompson.
23. Vuillemin, Jean. "Correct and Optimal Implementations of Recursion in a Simple Programming Language". *JCSS* 9 (1974), 332-354.
24. Ward, S. A. "Functional Domains of Applicative Languages". MAC TR-136, Project MAC, MIT, Sept. 1974.
26. Wadsworth, C. "Semantics and Pragmatics of the Lambda-Calculus" Ph.D. dissertation, Oxford(1971).
25. Wulf, W., et al. "HYDRA: The Kernel of a Multiprocessor Operating System". *CACM* 17,6 (June 1974), 337-345.